National Aeronautics
and Space Administration

Evaluating Elevated Temperature DRAM Stuck Bit Sensitivity at Room Temperature

Steven M. Guertin, Wilson P. Parker, Duc N. Nguyen, Sam W. Delaney, Paris Blaisdell-Pijuan, Joseph Vanacore, Gregory R. Allen, and Douglas J. Sheldon

Jet Propulsion Laboratory / California Institute of Technology

Pasadena, CA

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology,
Under contract with the National Aeronautics and Space Administration (NASA)
This work was funded by the NASA Electronic Parts and Packaging Program (NEPP)

To be presented by Steven M. Guertin at the 2019 IEEE Nuclear and Space Radiation Effects Conference (NSREC) July 8-12, 2019, San Antonio, TX

Stuck Bits at Launch

- Shortly after launch a mission experienced multiple stuck bits (Elpida 512 Mb SDRAM).
  - This was not a huge problem, but the mission did not have a plan for handling them (they did have EDAC, but stuck bits caused continuous error reports).
- While trying to understand the problem, we found a really frustrating problem.
  - Stuck bits persisted during mission for a "while".
  - But in the lab, if the temperature was raised "too much", most of the stucks rapidly disappeared.

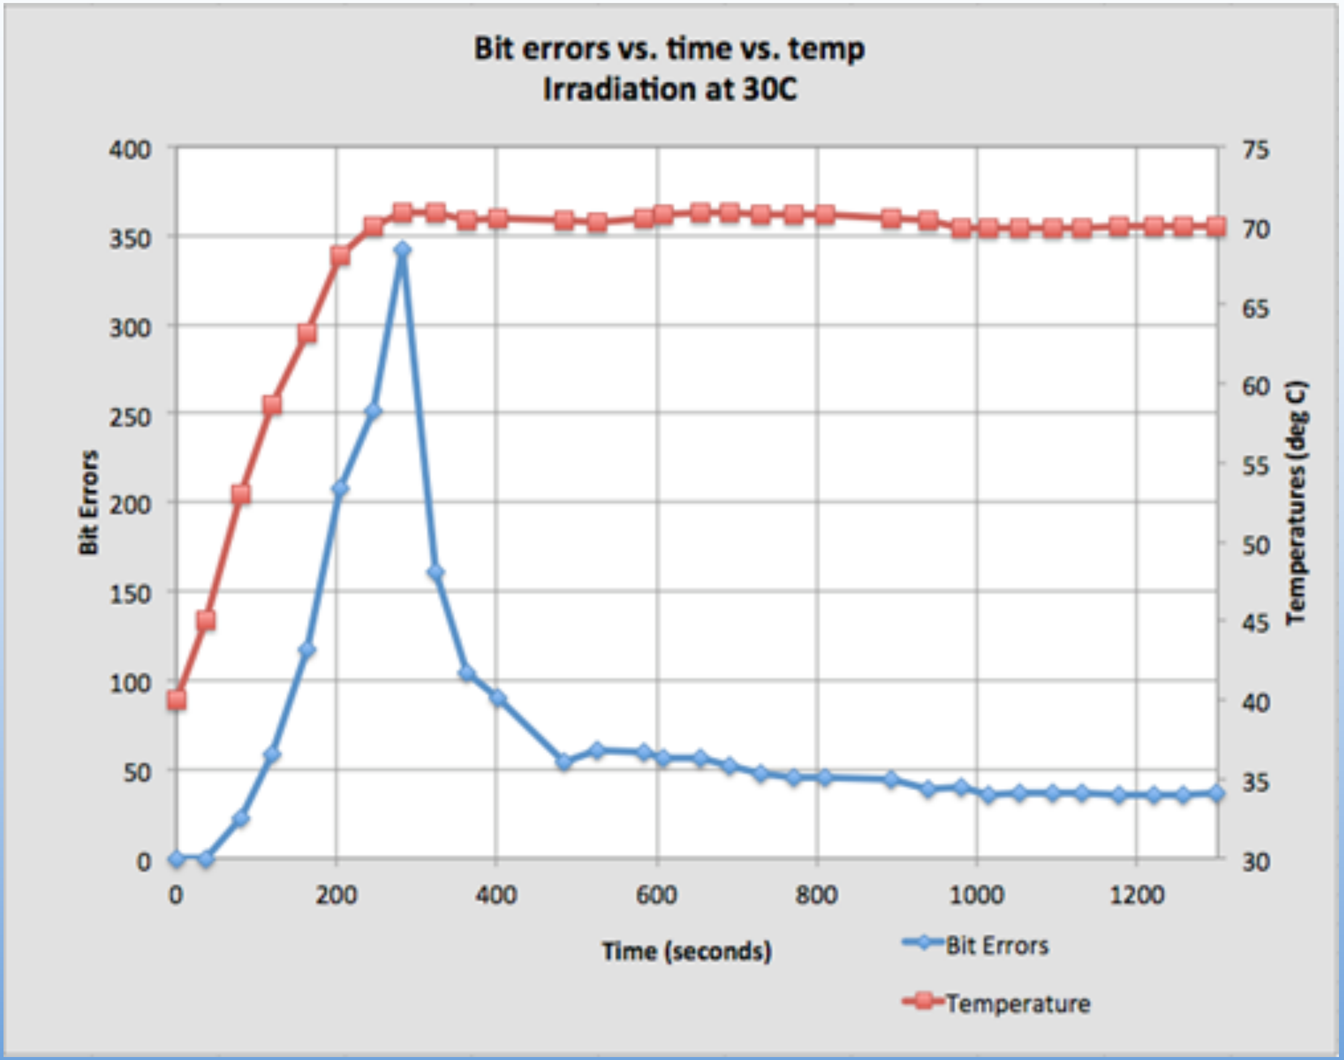

Recent Efforts

- Recent efforts to use a DDR2 (ISSI 1 Gb) device on a 300 krad(Si) mission have run into stuck bit issues.
- Initial testing for stuck bits revealed the possibility of 5000 stuck bits per part at end of mission at 75°C.
  - 2-4x annealing observed in first few weeks – at room temperature.
  - Limited additional annealing over months.
- At this rate, if a 9-device SECDED EDAC system is used, the probability of an uncorrectable error is >99%.
- This effort uses a method that does not involve heating the device
  - Clearly the large number of stuck bits requires accurate handling.
  - But we cannot underestimate the risk either.

Outline

- Background
- No-Heating Stuck Bit Method
- Is the Method Right/Good?
- Benefits
- Conclusion

National Aeronautics
and Space Administration

What is a Stuck Bit?

- It’s just a bit that can’t store data in the application.
  - Environment and operational parameters impact bits.
- Data stored in capacitor as charge decays over time
  - Refresh keeps data fresh
  - But if cell loses charge too fast, refresh will fail.

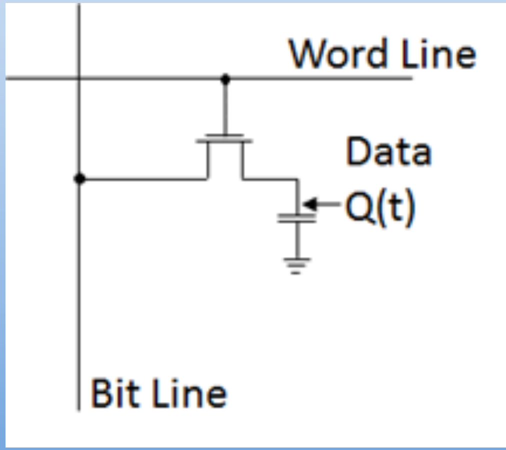

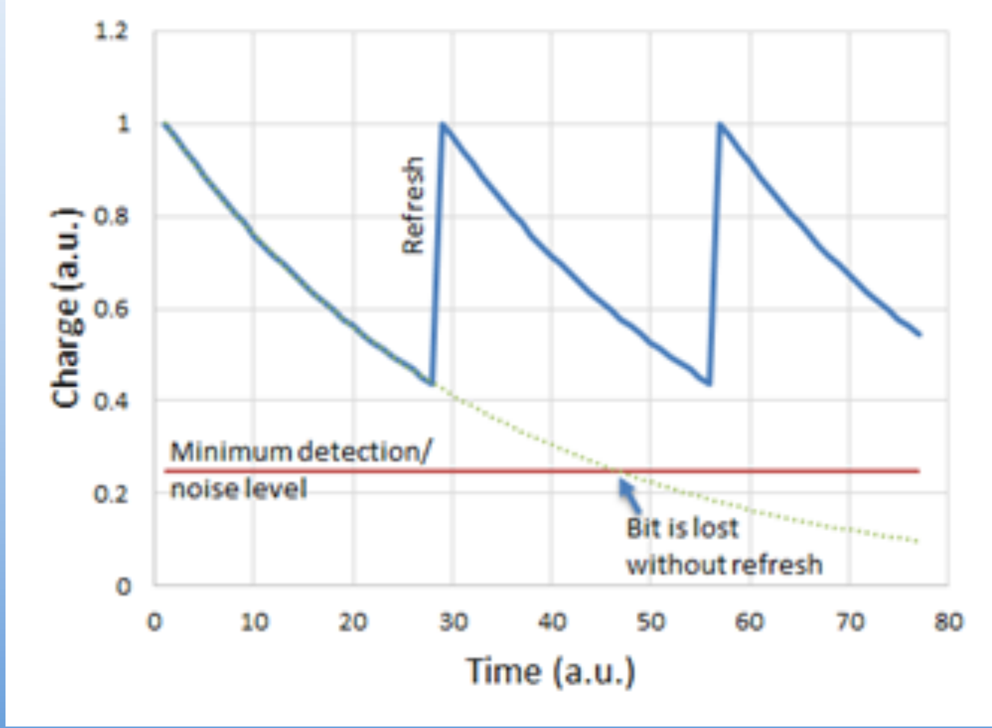

National Aeronautics and Space Administration

Measuring Retention Time

- The maximum time that a cell can retain data long enough to be refreshed reliably is the cell's retention time.
- Retention time studies have been used many times, including in manufacturer research on the performance of cells.

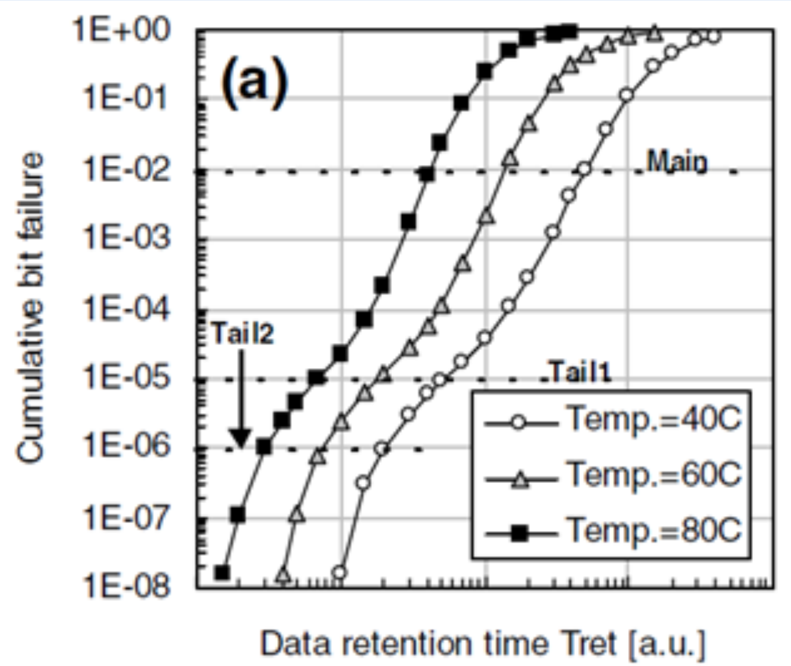

Wu, 2012

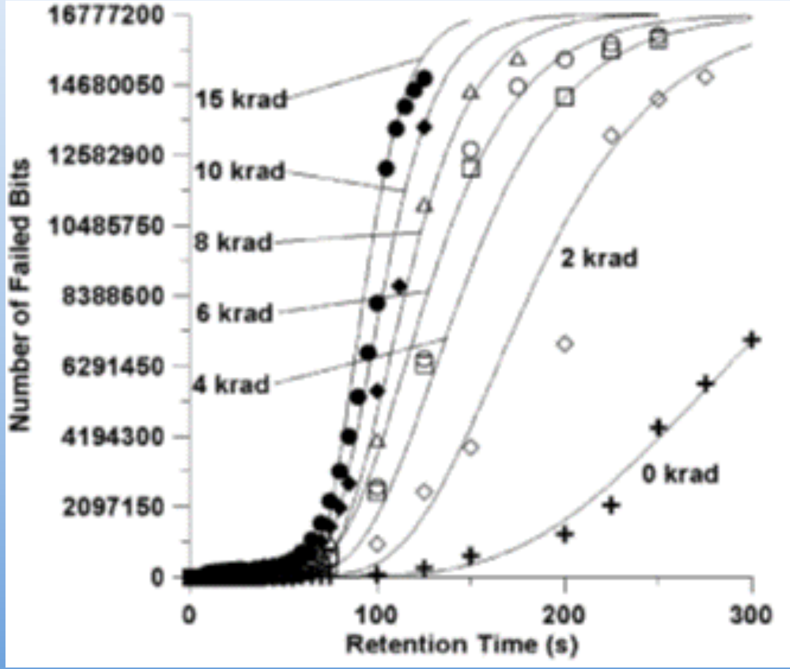

Scheick, 2000

Measuring Long Retention

- Manufacturers can isolate cells, but in application devices, retention of cells is measured by slowing down the refresh rate (even as long as days in some cases).
  - What does this mean for the device?
- Datasheet requirements typically give 64 ms as the minimum refresh rate (they don't specify a maximum).
  - They also often indicate 32 ms, if the device is used "hot".
  - The only reasons for this requirement might be: the cells may fail to store data (what we want to test), and failure to refresh the device may cause long-term damage (there are no indications of any issue).
    - We do not think the second reason is valid, but as radiation people, we can't make a blanket statement to the contrary.
  - So the keys are: make sure your data is reliable, and that there is no correlation of errors/behavior with amount of time spent doing long refreshes.
- Note that any issue with the interface to the device (because of lack of access) can be eliminated due to "self refresh" which allows the user to arbitrarily drop the device into a state where it waits for commands but can be essentially "hanging off the bus" for minutes to months.

How the Measurement Works

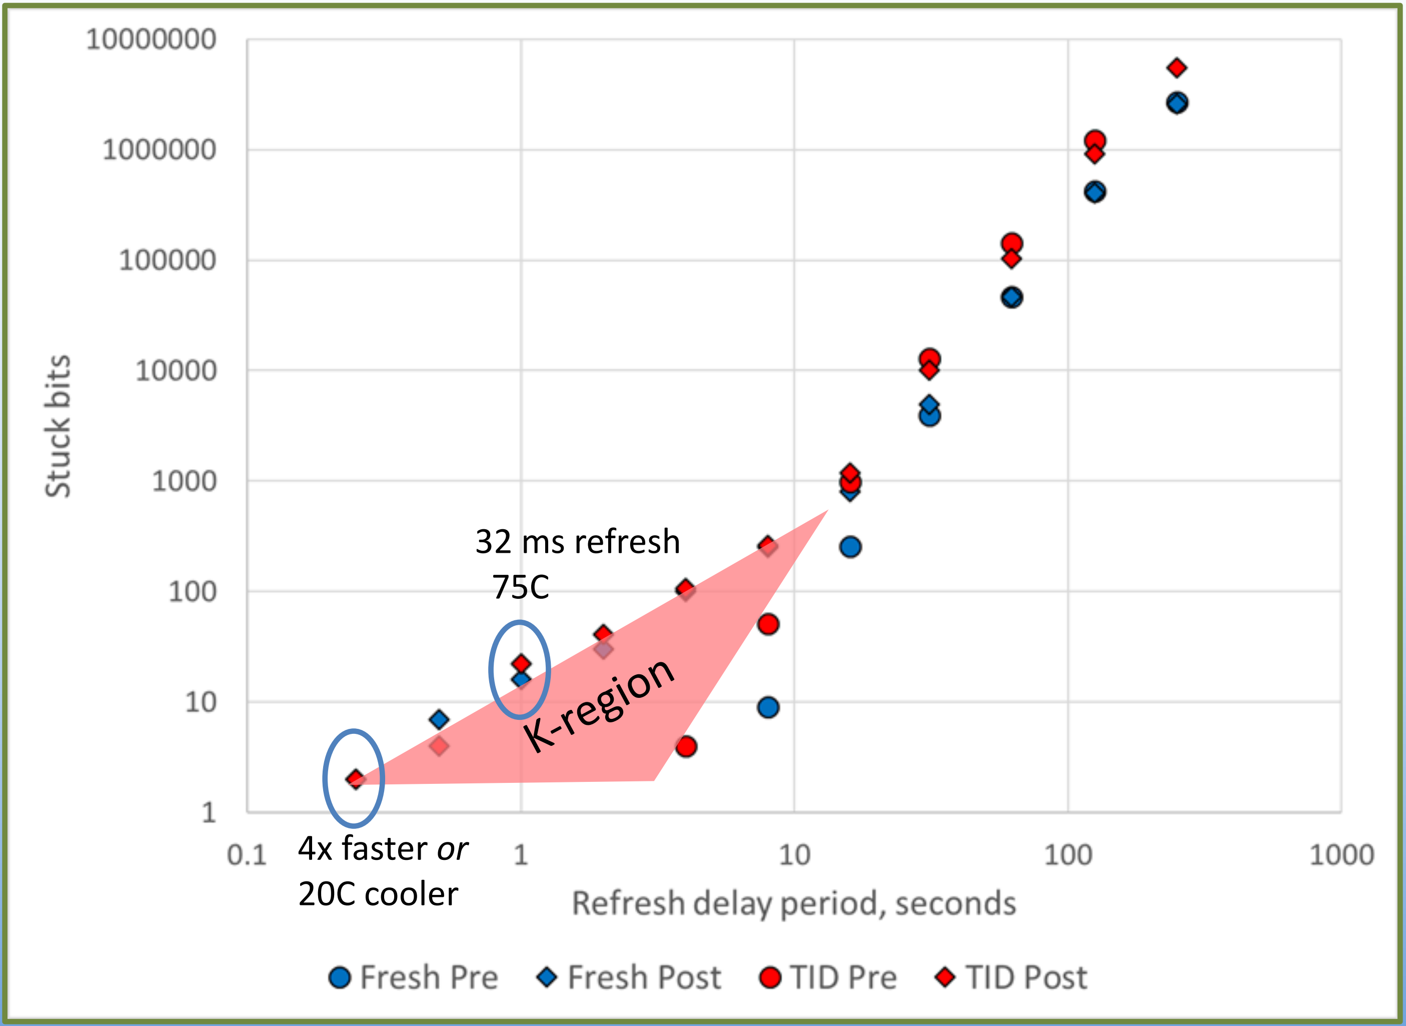

- The method is based on using a model for leakage current as a function of temperature.
- Then we increase the refresh period to make the amount of total leakage in the cells be about the same as if the cells were refreshed at 64 (32) ms, but at elevated temperature.
- Conceptually, the number of stuck bits moves “to the left” if the temperature decreases, or the refresh period decreases.
- Note that the “k-region” is an example of how different types of radiation can result in differences in the tail of the distribution – giving rise to more stuck bits.

Why is a Method Needed

- We can’t just heat the parts, because it destroys the data
  - In one of the earlier examples, the spacecraft held the parts in a 5-degree band between 50 and 55°C.
  - Lab testing proved very difficult because even a few degrees above 55°, all of the stuck bits annealed very quickly.
- The situation can’t even be worst-cased very well.
  - If you heat the parts to see the worst-case number of stucks…
- A model approach is required in order to predict the worst-case thermal performance without annealing.
- The main risk is that the method we’re working with is “off the datasheet”, which makes many engineers throw up red flags.
  - So the main issue, as indicated earlier, is whether or not anything is compromised.

Is the Method Right?

- Two primary questions:
  - Does this compromise the part or the value of the data?
  - Do predicted results actually contain the mission risk?
- Cells spend most of their time doing nothing
  - Datasheet only requires that cells be refreshed once every 64 ms, while they can be accessed as rapidly as every ~50 ns, giving a dynamic range of about a million – does not compromise the part.
  - The risk that the data is compromised is mitigated by directly verifying the temperature vs. refresh relationship.
- Mission risk is contained if a worst-case number of stuck bits can be estimated.
  - But it is possible the method may severely over-estimate risk… (Did that happen in the case of the 5000 stuck bits/device?).

Verifying the Model

- The two primary concerns are: Is the temperature to refresh model right? And, are the devices compromised.
- No impact to device performance by using long refresh periods has ever been observed.
- Temperature to refresh model may require fairly large systematic uncertainty:

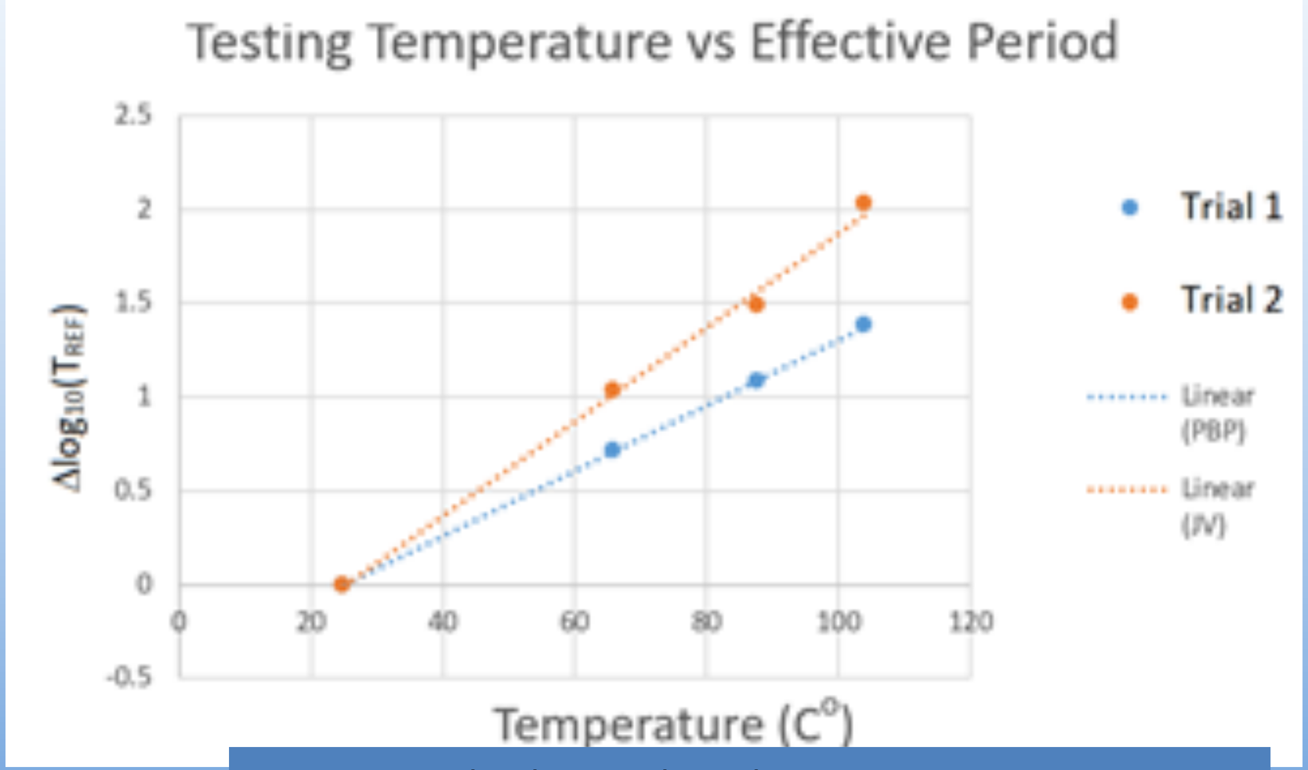

Exponential relationship shown to vary ~1.7 to 2.1 (2 means leakage current is doubled for each 10°C)

Ability to Build Datasets

- Testing at elevated temperature limits the validity of the dataset to the operating scenario chosen for testing.
- Users may need to optimize a user application, or need to know the evolution of stuck bits if multiple temperatures are used during a mission.
- The data collected using the model approach based on room-temperature evaluation, can be applied easily to other used cases.
  – Model dependency on minimized annealing, lower limit for temperature, and relationship between temperature and refresh period must be accounted.

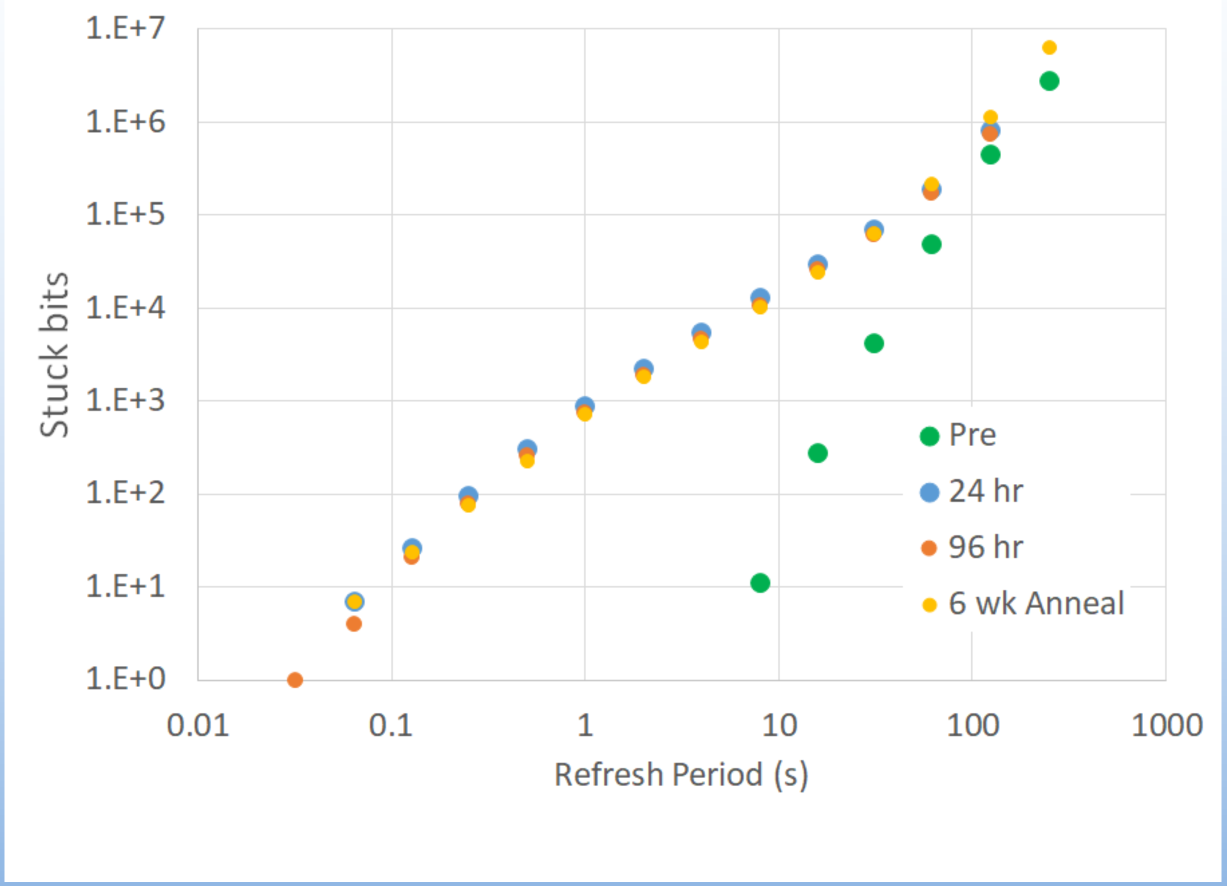

6 months later, this device still had ~700 stuck bits after adding 5 hours at 125°C, dropped to ~350 and after another 48 hours, dropped to ~50…

Conclusion

- Stuck bits are an issue for DRAM devices, and it is unclear how they will affect any given device.
- Stuck bit performance is worst when a device is operated hot, but if a device is evaluated at high temperature it can result in annealing of stuck bits and an under-estimate of lower-temperature performance.
- Time-performance of bits, evaluated by altering the refresh period, can be used to evaluate high temperature performance with a room-temperature evaluation.
  – Must take into account model uncertainty.
- Observations indicate deviation from datasheet only violates the cells' ability to store data, which is exactly what is being tested.
  – Periodic interaction with the device is not required (per Self-Refresh behavior).